

Appl. No. 09/614,489
Amtd. Dated July 25, 2005
Reply to Office Action of 4/27/2005

AMENDMENTS TO THE DRAWINGS

The attached sheets of drawings include 8 Replacement Sheets for Figures 1-8.

Fig. 2 has been amended to include the label "Port" on the hub from Port 9.

Fig. 3, Fig. 4, Fig. 5, Fig. 6, Fig. 7, and Fig. 8 have been amended to remove extraneous text.

No other changes have been made and no new matter is added. 8 Replacement Sheets and 7 Annotated Sheets are included herewith.

REMARKS

Claims 29-55 are currently pending. Claims 29-31, 36-43, 46-51 and 53-55 are rejected under 35 U.S.C. 102(e), and claims 32-35, 44, 45 and 52 are rejected under 35 U.S.C. 103(a). The Applicant traverses the rejections and makes detail explanation to distinguish the difference between the present application and the reference.

There are no amendments - Applicant respectfully requests entry of the arguments to place the application in better condition for appeal if not allowed in its entirety.

In "Response to Arguments" on page 10 of the Office Action, the Office states that the description in Tosey, "connecting network computing device can assume the Ethernet port of the network peer device is in operation", disclosed or otherwise defines a "managed network switch". The Applicant respectfully disagrees and provides further explanation herein and in the accompanying Declaration. This incorrect assumption is the basis for the further rejections and once traversed, the rejections relying upon this basis are also traversed.

This description in Tosey never disclosed the same or similar function of a Managed Ethernet Switch, which is defined in present specification and distinguished from the Tosey implementation as follows:

An 'Ethernet Switch' or 'Layer 2 switch' is a device that transmits message packets unchanged from one of its ports to another, using rules that are dependent only upon the destination MAC address of the message. Such devices are becoming the preferred interconnection devices for large Ethernet networks, since they do not require significant configuration. This is as opposed to 'routers', otherwise known as 'Layer 3 switches'.

(Present Application; Page 2 – beginning on line 28)

A ‘Managed Ethernet Switch’ is an Ethernet switch which includes a management entity conforming to the reporting requirements of RFC 1493, and which therefore specifically may be interrogated to determine which port of the device was used recently to receive a message from a particular MAC address.

(Present Application; Page 3 – beginning on line 5)

The Office refers to Tosey Col. 7, lines 6-33 regarding the use of Ethernet. However the use of Ethernet as a network does not anticipate or render obvious the methodology of the present invention in conjunction with a Managed Network Switch and the processing therein to identify a location of a failed target device to isolate a single known IP Address for that target device. An Ethernet hub is not a Managed Ethernet Switch – it is still a hub. And, the Tosey methodology clearly indicates that it is not intended to be used with a Managed Ethernet Switch.

More specifically, Tosey section in Col. 7, lines 6-33 recites:

In yet another technique, a network device can attempt to access a resource controlled by the peer device to test the connection. This method attempts to access a TCP or UDP port on another device, which may be denied. When a peer device denies access, or sends a "connection refused" response, the connecting network computing device can assume the ethernet port of the network peer device is in operation.

Returning to FIG. 4A, after the computing device executes the link test, the logic proceeds to block 108 where the network computing device 21 determines if a peer network device has returned a response. If the network computing device 21 does not receive a response, the process continues to block 106 where it executes a wait state for T_1 seconds. For block 106, a wait period, T_1 , of 5 to 10 seconds is sufficient. However, if a peer network device returns a response at block 108, the process continues to block 110 where the network computing device 21 records the IP address of the responding peer computing device in one of its memory devices. In this section of the process, it is only necessary to record the IP address of the first response. However, it may provide additional security if additional IP addresses of other responding peer network computing devices are recorded. At this time, the network computing device may also record the network response time of the link test for each corresponding node. This value may be used for the wait state in block 114, as described below.

Tosey does not describe any Managed Ethernet Switch, and this section actually describes the processing of Tosey to continuously ping the peer network devices to establish a list of IP addresses of all responding peer devices. The network computing device 21 stores the IP address of the responding peer computing device. This does not match a MAC address to the IP address nor would it identify the port location of the failed target device. Tosey fails to disclose or describe a system that functions as claimed. And, it is improper for the Office to impart elements clearly

Therefore, the claimed elements that incorporate a “Managed Ethernet Switch” refers to a specific network switch as defined and explained in the specification. And, the processing in conjunction with the Managed Ethernet Switch distinguishes it in the manner of operation as a standard Ethernet hub cannot perform in the same manner. Applicant also refers the Office to the Rule 132 Declaration for further details distinguishing a Managed Ethernet Switch as compared to hubs and routers – whether Ethernet or some other network type – and to the manner of operation of such switches.

An Ethernet switch automatically divides the network into multiple segments, acts as a high-speed, selective bridge between the segments, and supports simultaneous connections of multiple pairs of devices which don't compete with other pairs of devices for network bandwidth. It accomplishes this by maintaining a table of each destination address and its port. When the switch receives a packet, it reads the destination address from the header information in the packet, establishes a temporary connection between the source and destination ports, sends the packet on its way, and then terminates the connection. But when a hub receives a packet of data (a frame in Ethernet lingo) at one of its ports from a device on the network, it transmits (repeats) the packet to all of its ports and, thus, to all of the other devices on the network. If two or

more devices on the network try to send packets at the same time a collision is said to occur. When that happens all of the devices have to go through a routine to resolve the conflict.

The present invention employs the Managed Ethernet Switch to associate a MAC address with a port location of a target device so that a single IP Address can be associated to an individual target device. Tosey does not describe such a methodology or functionality and instead details the processing for a standard hub.

The Office also states in the "Response to Arguments" that Tosey discloses a method and system for receiving from a target device on the network a request to be assigned an IP Address. However, this does not identify a single known IP address for a failed target device according to the present claims and it does not locate the port associated with the device. As noted in Claim 29 of the present application - "the request including a Media Access Control (MAC) address associated with the target device." Having a list of IP addresses for target devices does not identify which of those IP Addresses is the correct one for a failed target device on a port. Tosey does not describe or infer any scheme to identify a single IP Address for a failed target device by locating a physical location of the target device by the steps noted in the present claims. The attached Rule 132 Declaration provides additional explanation and clarification

Claims Rejections - 35 USC §102(e)

The Office rejected claims 29-31, 36-43, 46-51 and 53-55 are rejected under 35 U.S.C. 102(e) as being anticipated by Tosey et al., U.S. Pat. No. 6,392,990. A rejection based on anticipation requires that a single reference teach every element of the claim (MPEP § 2131). "The identical invention must

be shown in as complete detail as is contained in the ... claim." *Richardson v. Suzuki Motor Co.*, 9 USPQ2d 1913, 1920 (Fed. Cir. 1989). Or stated in another way, a "claim is anticipated only if each and every element as set forth in the claim is found, . . . described in a single prior art reference." *Verdegaal Bros. v. Union Oil Co. of California*, 814 F.2d 628, 631, 2 USPQ2d 1051, 1053 (Fed. Cir. 1987).

In the rejection of claims 29, 46, 48 and 53, Examiner cited the "hub" in the Tosey reference in order to anticipate the "Managed Ethernet Switch" of the claims in the present application. As known by one of the ordinary skill in the art, "hub" and "Managed Ethernet Switch" are two very different devices. Although hubs and switches both "glue" devices in a network together, they have totally dissimilar characteristic in operating principles and functions. They are not equivalent in form or function. For example, a hub does not respond to a query having a MAC address and return a port number suitable to identify the IP Address for the replacement device as described and claimed in the present application. There is no MAC address processing according to the claimed invention. While Tosey uses the term "MAC address" in Col 8, lines 44-58 – it is used as shown in Tosey Fig 5A/5B with respect to updating the new IP address of the network device. This is not remotely close to the claimed methodology.

Applicant has submitted a Rule 132 Declaration and respectfully requests that the Office review the Declaration for a detailed explanation of various network devices. The Declaration also explains and clarifies some of the distinguishing attributes between Tosey and the present invention. A Declaration was not submitted previously as it do not seem necessary to describe prior art switches, however upon review of the Final Rejection, the Inventor believes that such a Declaration will expedite processing by carefully explaining the state of the art at the time of the invention to distinguish Tosey.

The Office refers to Tosey Figure 3 (see below) wherein network computing devices 31, 32 have paired/redundant Network Interface Cards 37, 38 and 39, 40 that are coupled to hubs 33, 34. The aim of the Tosey invention is to give the appearance of continuous fault-free operation of the network computing device, by having the substitution of the network interface cards be automatic. In Tosey, the component to be replaced is the Network Interface Cards (NIC), and the replacement was necessarily anticipated by the designer of the network computing device – having redundant NIC cards in the network computing device. The network computing device itself maintains all knowledge of the required IP addressing and other configuration details. When the network computing device 31, 32 determines that one of its redundant Network Interface Cards 37, 38 and 39, 40 is unable to continue operation, the network computing device 31, 32 performs an internal reorganization whose effect is to substitute one of the spare communication interfaces, and reconfigure this spare with the known IP address parameters of the communication interface which it is replacing. However, knowing the IP address of a failed device does not identify the physical location of a port on a hub. If, for example, a technician knew the IP address of a failed device, and there were 6 devices hooked to the ports of a hub – Tosey does not identify a port on the hub for the failed device. Furthermore, as described in the present application, there are situations wherein it would not be possible to isolate a single IP address.

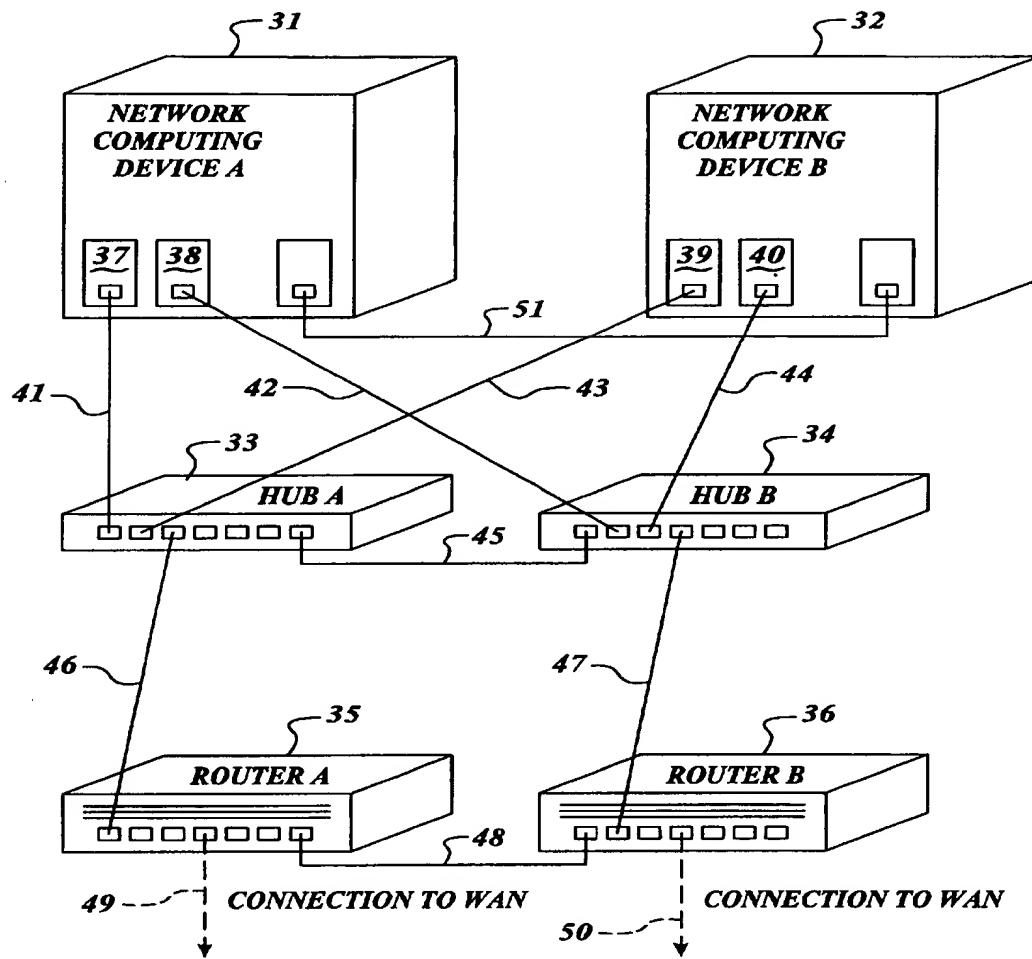
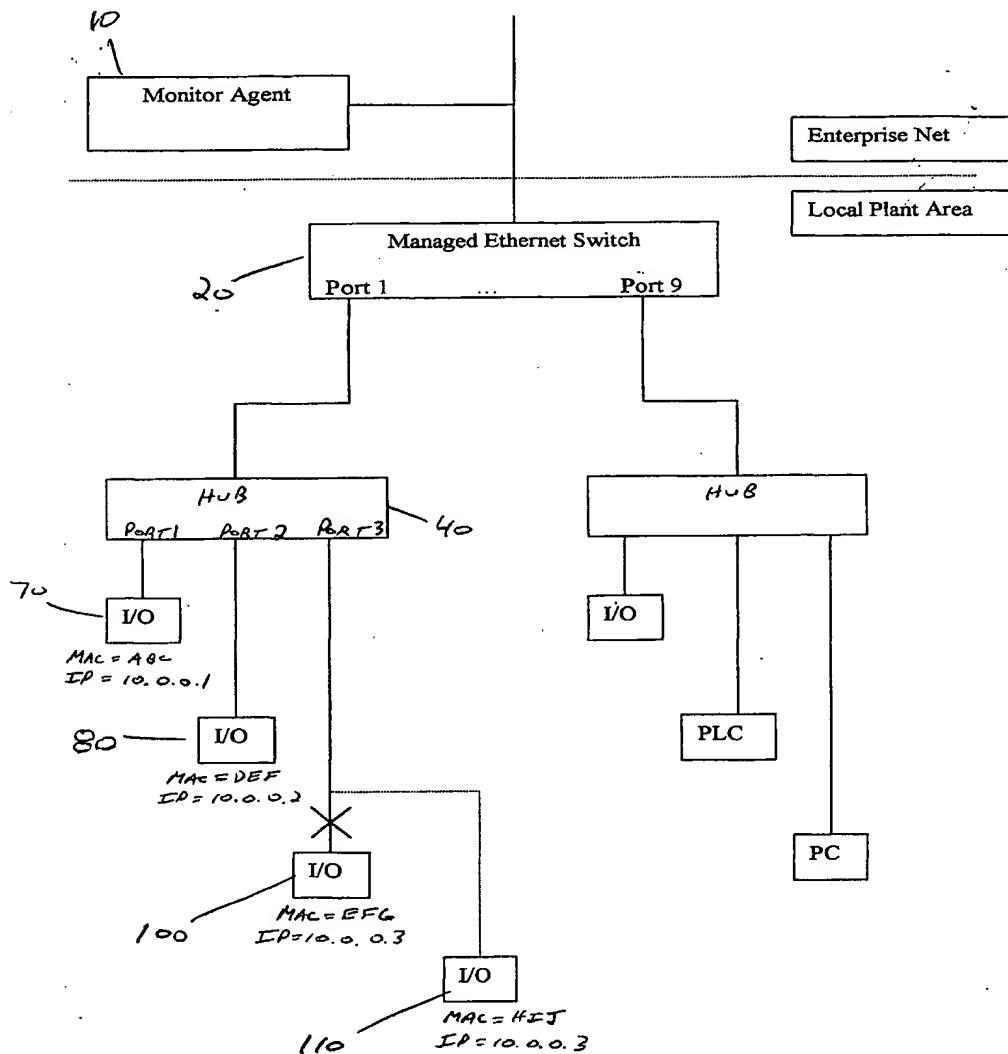


Fig. 3 ↗
30

This is in sharp contrast to the present invention that employs a Managed Network Switch and identifies the failed target device location. The following Figure 2 of the present application illustrates the distinguishing attributes from a system perspective. There is a Managed Ethernet Switch 20 that is coupled to hubs 40 which in turn couple to target devices 70, 80, 100, 110 with different IP Addresses as shown. The processing as claimed allows the present invention to replace the non-functioning device 100 and have the system assign the proper IP Address according to the claimed method – wherein the hub 40 has multiple

possible IP addresses. With respect to Tosey, the structure is different, the elements are different and the processing is different as further detailed herein and in the attached Declaration.



Claim 29 accurately reflects the methodology of one embodiment of the present invention and thus distinguishes the present invention. Claim 29 recites:

A method for determining the correct Internet Protocol (IP) address for network-connected devices, comprising:

receiving from a target device on the network a request to be assigned an IP address, the request including a Media Access Control (MAC) address associated with the target device;
issuing a query to one or more managed Ethernet switches on the network, each switch having a number of ports, where each query specifies the MAC address and requests that the queried managed Ethernet switch report the number of any port on which was received data sent by a device having the specified MAC address;
receiving replies to one or more of the queries; and
in response to determining that one of the queried managed Ethernet switches and a port number reported by that switch corresponds to a single known IP address, assigning that known IP address to the target device.

Tosey does not receive a MAC address and IP address from a target device; Tosey does not communicate with a Managed Ethernet Switch to obtain port information via a query based on a MAC address; and does not process the information to identify a single IP address for the target device. For at least these reasons, reconsideration and allowance is respectfully requested for all claims rejected under 35 USC 102.

The present invention issues queries to the Managed Ethernet Switch to obtain the number of the port on which was received data sent by a device having the specified MAC address from the target device. A hub does not provide such information. For further clarification, the definition of "Ethernet switch" and "Hub" in the technology dictionary is repeated below for reference. While the CAFC has indicated a preference to the specification definition, the dictionary definition clearly supports the Applicant's position.

Hargrave's Communications Dictionary (Frank Hargrave, copyright 2001 by the Institute of Electrical and Electronics Engineers, Inc. (IEEE), published by Wiley)

Ethernet switch

A data switch that allows users on different network segments to exchange data. When users on different segments exchange data, an Ethernet switch dynamically connects the two separate Ethernet channels without interfering with other network segments. The switch can create multiple independent connections between separate segments, allowing multiple parallel data exchanges. This multiplies network bandwidth without modification to Ethernet end station hardware or software.

hub

A network component that:

- Serves as a common termination and distribution point for multiple nodes (a physical star topology) and
- Accepts a signal from one node and redistributes it to other nodes.

Since the explanation of the differences between “Hub” and “Ethernet Switch” unequivocally supports the proposition that the hub in the Tosey reference cannot be applied to the present application, the rejections based on the hub to anticipate the managed Ethernet switch are inappropriate and should be reconsidered.

The Office also rejects Claim 30 as being anticipated by Tosey. Although Tosey is maintaining a database, this database performs a totally different function from the database in the present invention. In Tosey, a database is maintained by the local network computing device, to record the identity of the local interface which will be used to send and receive messages to each of a number of peer computing devices, identified by IP address. So at a later time, when apparent failure to communicate with one or more IP addresses is detected, the local network computing device can, by consulting the database, distinguish an apparent failure of its local interface from a (real) failure of one or more peer devices. By doing this, Tosey avoids making (futile) substitution of a network interface if at least one of the targets normally ‘seen’ through the interface continue to be seen.

In the present invention, the database is maintained in the BOOTP server device, and specifically associates an IP address with a combination of a managed switch IP address and a reported port number on such managed switch from which messages emitted from the IP address have been recorded.

Tosey makes no mention of the use of managed switches, nor the port numbers of such switch which are reported by the managed switch as being attached to a device with a particular IP or MAC address.

The Office also rejects Claim 31 as anticipated by Tosey, wherein Tosey is concerned with detection of a failure of a network interface, through non-response to a stimulus of some type sent across the network, and then adjustments of a database as a consequence of substitution of that interface.

Claim 31 of the present invention refers to determining, using the database described in claim 30, that a newly discovered device, requesting assignment of an IP address, is in fact already known to the server by having the port number of the managed switch to which it is connected recorded in a row of the database. Tosey makes no mention of managed switches and port numbers thereof.

The Claim 36 anticipation rejection is in view of Tosey. Tosey describes use of a variety of network stimulus techniques from which it may be deduced, and from the non-response to such stimulus, that a network fault of some type exists. The types of test mentioned in Tosey include ICMP PING request, ARP request, TCP or UDP connections. Each of these is a widely used technique and well known. All of the techniques are described and encouraged in the Internet standard RFC documents available prior to 1990.

Claim 36 methodology includes the identification of a physical wiring location determined using a query to one or more managed switches, purely as a further method of reducing ambiguity, and allowing the present invention to be applied to networks involving some unmanaged components in addition to the managed switches fundamental to the mechanism. This refinement reduces the cost and increases the attractiveness of the present invention. Tosey does not mention use of managed switches and port numbers obtained therefrom.

Claims 37, 38 and 38 are also rejected as being anticipated by Tosey. These claims elaborate some of the network stimulus techniques which maybe used to determine the apparent failure of a device which, in conjunction with the location determination described in claim 29, result in determination of the IP address of the replacement device. One of the improvements obtained by associating apparent failure of a device with the replacement IP address determination are to eliminate accidental issuance of an IP address which conflicts with one already in use, and to extend the applicability of the invention to networks involving unmanaged (and therefore lower cost) devices using the refinement in claim 36

With respect to claim 41, also rejected as being anticipated by Tosey, claim 41 relates to the status of a set of devices recorded in a database to resolve ambiguity in the case where a replacement device requesting IP address assignment has been found already to be on a port of a managed switch to which is attached more than one potential target device. Tosey makes no reference to use of port numbers from managed switches.

The anticipatory rejection of claim 43 is traversed as Tosey does mention sending out messages to inform other devices of a change in IP address assignment. The IP assignment which is being reported in Tosey is that of the sending computer itself, as a result of the reorganization of its network

interfaces prompted by the apparent network device failure. The present invention sends out messages advising a target device, which has previously requested allocation of address using the BOOTP protocol, of the successful satisfaction of that request. The IP address being carried is the address of the target of the message, not that of the source.

The rejection of claims 46 and 47 are also traversed as Tosey makes no reference to managed switch or port number obtained therefrom.

In summary, a “hub” is different from an “Ethernet switch”, and cannot provide the port information using the target device MAC and IP address. Therefore Tosey fails to disclose all elements of the independent claims 29, 46, 48 and 53, so the Applicant respectfully requests the Examiner to withdraw the rejections. Claims 30-31 and 36-43 depend on the independent claim 29, claims 47 depends on the independent claim 46, claims 49-51 depend on the independent claim 48, and claims 54-55 depend on the independent claim 53, so reconsideration of rejection of these claims is also respectfully requested.

Claim Rejections – 35 USC § 103

The Office has quoted the statute from 35 U.S.C. 103(a), which is referenced herein. The Office rejected claims 32-35, 44, 45 and 52 as being unpatentable over Tosey et al. The Applicant traverses the rejections and respectfully includes a detailed explanation to distinguish the present application from the Tosey reference.

According to the MPEP §2143.01, “[o]bviousness can only be established by combining or modifying the teachings of the prior art to produce the claimed invention where there is some teaching, suggestion, or motivation to do so found

in either the references themselves or in the knowledge generally available to one of ordinary skill in the art."

Once again, the problem being solved by the Tosey invention is a very different one from that solved by the present invention and therefore the solution of Tosey is distinguishable.

The elements in the independent claims of the present invention includes the interrogation of a Managed Ethernet Switch on the network, where a managed switch is one which will report the identity of the incoming network port on which was previously received a message from a particular MAC address. Tosey makes no mention or inference of issuing a query to a managed switch in order to determine the port to which is attached a device having a particular MAC address.

In contrast, the only devices mentioned in the Tosey invention include Network Computing device, LAN, WAN, Router, Hub, and Network interface card. None of these are similar to a managed switch nor do they provide the port information which is processed according to the teaching of the present invention.

Furthermore, merely placing a Managed Network Switch in Tosey does not produce the present invention as Tosey does not give any indication of using the target device MAC address and the port number to derive the IP address as claimed in the present invention.

A ‘Managed Ethernet Switch’ as described in the present invention operates at the ‘media access control’ layer and uses MAC address information to make its switching decisions. The present application specifically defines the

‘managed switch’ in the patent specification so that it is clearly differentiated from the other networking devices.

For at least the reasons set forth herein and supported by the Rule 132 Declaration, a “hub” cannot be equated to the “Ethernet switch” of the present application, as it does not support the functions of a Managed Ethernet Switch. “Hub” and “Ethernet switch” have different operating principles and functions, and there is no teaching, suggestion, or motivation to substitute “Ethernet switch” with “hub” found in either Tosey or in the knowledge generally available to one of ordinary skill in the art. Contrarily, people of ordinary skill in the art understand that “hub” and “Ethernet switch” are two different instruments with totally dissimilar characteristic in operating principles and functions.

As argued herein and supported by the Declaration, the rejections of the claims under 35 USC 103 is traversed as Tosey makes no mention of managed switch and port number which are aspects to the mechanism of the presently claimed invention. The present claims use a managed switch and issues a query to the managed switch to find the port number to which is attached a device with a particular MAC address. This is not at all obvious, especially since RFC 1493 was issued in the year 1993, and yet no mechanism for automatic IP address assignment using the information in this specification has yet been produced.

According to the MPEP §2141.01(a), “In order to rely on a reference as a basis for rejection of an applicant’s invention, the reference must either be in the field of applicant’s endeavor or, if not, then be reasonably pertinent to the particular problem with which the inventor was concerned.” (citation omitted). And the example of “Analogy in the Electrical Arts”, the same section of the MPEP also explains the situation of inappropriate rejection based on obviousness.

Furthermore, according to the MPEP §2143.03, If an independent claim is non-obvious under 35 U.S.C.103, then any claim depending therefrom is non-obvious. Claims 32-35, 44 and 45 are dependent on claim 29 and should be allowed as claim 29 is non-obvious. Reconsideration of the rejection of claims 32-35, 44 and 45 is therefore respectfully requested.

Telephone Interview

Present Office policy places great emphasis on telephone interviews initiated by the examiner. For this reason, it is not necessary for an attorney to request a telephone interview. Examiners are not required to note or acknowledge requests for telephone calls or state reasons why such proposed telephone interviews would not be considered effective to advance prosecution. However, it is desirable for an attorney to call the examiner if the attorney feels the call will be beneficial to advance prosecution of the application. MPEP§408

Applicant believes the above remarks to be fully responsive to the Office Action, thereby placing this application in condition for allowance. Applicant requests speedy reconsideration, and further requests that Examiner contact its attorney by telephone, facsimile, or email for quickest resolution, if there are any remaining issues.

Respectfully submitted,



Vernon C. Maine, Reg. No. 37,389
Scott J. Asmus, Reg. No. 42,269
Andrew P. Cernota, Reg. No. 52,711
Attorneys/Agents for Applicant

Cus. No. 24222
Maine & Asmus
PO Box 3445
Nashua, NH 03061-3445
Tel. No. (603) 886-6100, Fax. No. (603) 886-4796
Info@maineandasmus.com



IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of: Swales, A.

Group Art Unit: 2151

Serial No. 09/614,489

Examiner: DINH, K. Q.

Filed: 07/11/2000

Dkt No: LAN01

For: AUTOMATIC DETERMINATION OF CORRECT IP ADDRESS
FOR NETWORK-CONNECTED DEVICES

To: Mail Stop AF
Commissioner for Patents
PO Box 1450
Alexandria, VA 22313-1450

From:

24222

Dear Honorable Commissioner:

This declaration is offered in support of the above application for patent.

RULE 132 DECLARATION OF ANDREW G. SWALES (37 CFR 1.132)

I am presently the President of Scorpion Controls, a successful New Hampshire business that provides consulting services to many companies throughout the United States. I am an independent contractor in the automated networking field and regularly design, test, develop and integrate automated manufacturing systems employing networked devices.

I have extensive experience in the fields of networking and computer systems, including Ethernet, TCP/IP routing, BOOTP address assignment, and SNMP monitoring, primarily obtained during my work as Chief System Architect for Modicon Corp.

I am a member of a number of professional organizations including the Institute for Electrical and Electronics Engineers (Member since 1982)

I graduated from the University of Cambridge, England with a Master of Arts degree in Physics and Computer Science, having achieved highest honors in both parts in 1973.

I began working on networked systems in 1976 while employed by Kent Automation, now part of Brown Boveri Group, a European manufacturer of computer systems for industrial control.

From 1981 -1999, I acted as System Architect, and became Chief System Architect for Modicon Corporation in Andover, Mass. (now part of Schneider Electric). In this role I was responsible for defining computer system and network architectures for a range of Programmable Controller and telemetry products, including the network protocols familiar to most in the field known as Modbus, Modbus Plus, Modbus II, and finally Modbus/TCP.

During that period, I also represented the company and attended meetings of the IEEE standards committee defining the computer network known as IEEE 802.4, one of the early alternatives to Ethernet in the industrial control arena.

I developed and published the specification for the Modbus/TCP protocol, which was adopted as a standard by all major industrial networking companies.

In 1999, I left Schneider Electric but was retained by my previous employer on long term consulting assignment for my Network Protocol expertise, supporting customers such as IBM Semiconductor and United Parcel Service. One of my major projects during this time was to define a set of Conformance Tests for the Modbus/TCP protocol which IBM would use when qualifying vendors of networking equipment

My current major project is in the field of Simple Network Management Protocol (SNMP) software, where I am assisting in the design and development of a software package allowing the layout of a computer network to be ‘mapped’ automatically, using techniques similar to those used in the invention for which this patent has been applied.

I am an inventor on several patents in the networking field including U.S. Pat. No.’s 6,151,625; 6,233,626; 6,321,272; 6,434,157; 6,466,995; 6,640,140; and 6,760,782. While these inventions are not relevant to the present invention, I was intimately involved in the patent process with these

patents and have a fair understanding of the process and requirements. Thus, while not a patent professional, I have a basic understanding of the requirements and the importance of maintaining integrity in the patent system.

Tosey Patent U.S. Pat. No. 6,392,990

The Tosey patent describes an elegant method of solving a fundamentally different problem than my invention described in the patent application. In Tosey, the objective is to allow a combination of computers, cables, hubs and routers to recover without human intervention from a single point failure of a single cable, interface, or hub.

This is arranged by having multiple ('redundant') network interfaces, each of which is connected to a port of a hub, and connecting the hubs together, in such a way that there is more than one independent communication path from any one such computer to another. Any single fault of a cable, interface, or hub, therefore, although it may prevent communication along the original path, will leave an alternative path for communication between the two parties.

As an important enhancement, the computer is programmed to perform a series of continuous tests on its ability to communicate to its peers. If it is determined that there is a problem with communication, the network is checked to see if it is possible to reconfigure a standby interface to take on the identity of the interface which is no longer accessible. If so, such reconfiguration of the interfaces is performed.

As a result, the standby interface is able to receive messages originally intended for the failed interface, and there is no need for the peer computer to adjust anything in order for communication to be restored. The advantage of doing all of this is to give the illusion to other computers and routers that no failure had occurred, or that the duration of the failure was very brief.

A key assumption, however, is that the computer concerned is already aware of its 'role' on the network, and which IP address is supposed to be allocated to each interface.

It is necessary that each computer is aware - ahead of time - which IP addresses it is supposed to use, because in the failure condition it is going to need to perform reorganization of its network interfaces without relying on accessibility to an external server of any kind. After all, the failure it is trying to overcome is likely to impede communication to such server.

Tosey also makes allowance for the possibility that a network interface for an individual computer may have failed, and that it is necessary to perform manual removal and replacement of the component. The same software logic which would be used to determine the correct IP address to associate with the standby interface in the case of failure can also be used to determine the IP address of the newly replaced interface.

But this capability is already present in almost all computer systems with replaceable interfaces, such as the common IBM PC on everyone's desktop. When a network interface card is replaced, the software would automatically assign to the new card the same IP address as was previously assigned to the old card. So the value of this capability in practice is very limited.

Tosey makes brief reference to the initial assignment of all IP addresses on a computer either being made by static assignment, or by a dynamic technique such as DHCP or BOOTP. Such techniques were widely known and used by all network designers, and the use of BOOTP or DHCP is not fundamental to the Tosey invention. Thus, the technology used in the present invention was available at the time of Tosey and the present invention is not merely an extension of Tosey using newer technologies.

Present Invention

By contrast, the present invention is designed to solve the following problems:

An industrial control system, consisting of many parts connected together by a network, and there is a failure of a part.

It is understood that in some instances the solutions to two different problems may involve similar schemes – however such is not the case between the Tosey invention and the present invention.

It may be helpful to the understanding of the present invention to paint a picture of a typical situation that prompted me to arrive at the present invention. The customer in an industrial facility typically holds various replacement parts in a stockroom or similar location, where the parts are virtually indistinguishable. Upon the occurrence of a failure of one of the parts, an electrician is dispatched to substitute the failed device, which he typically does by removing all electrical and network cables from the device, then replacing the physical device, and finally replacing the electrical and network cables.

In systems lacking the capability of the present invention, the part which has just been replaced needs to have its communication parameters explicitly configured before communication can resume. A variety of techniques are in common use to do this, all are cumbersome and error-prone. The present patent application describes some of the techniques in the background section. All of them require either specialized knowledge on behalf of the electrician, or coordination with a network specialist of some sort. The end result in all cases is increased cost and downtime while the part is replaced.

But the present invention provides a solution not previously accomplished. Almost all modern network devices are capable of getting their IP address and other parameters from a BOOTP service elsewhere on the network. The present invention takes such a standard BOOTP server and enhances it by adding an automatic recognition feature based on information from one of two sources:

1. The 'port number' reported as being directed towards a particular network device, by knowing only its MAC address. This feature is available when using managed Ethernet switches from almost all vendors, and specifically all which claim to support the 'SNMP Bridge MIB' defined in RFC1493
2. The current presence or absence of network devices as seen by a conventional 'ping scanner'. Such a scanner is ordinarily used to provide advance warning to network professionals of

devices being added or removed from a network. But in the case of the present invention, it is used to resolve ambiguities in the possible IP addresses which the BOOTP server would respond with.

It is this combination of BOOTP service and port number recognition which is one of the important characteristics of the present invention.

From previous comments in office actions on this application, some clarifications about the fundamental capabilities of a ‘Managed Ethernet Switch’, and how these capabilities are employed in the present invention should help distinguish the present invention. It should be understood that nothing in the Tosey patent makes any reference to Managed Ethernet Switches of their capabilities. By contrast, Tosey only refers to routers and hubs, devices which have very different roles in the networking process.

It is necessary to understand a little more closely the intended functions of Bridges and Switches, and how they differ from hubs and routers, in order to recognize the significance of the present invention.

Network Bridges

The predominant technology used in local area networking today is Ethernet, based on developments by Digital Equipment Corp, Intel Corp, and Xerox Corp in the late 1970’s.

This networking method originally required that all information be sent in the form of ‘packets’ or ‘datagrams’, which are transmitted on a cable to which all stations are equally attached. Such an arrangement is known as a ‘broadcast bus’.

Because the channel is shared, it is necessary that all transmitters hesitate before sending a message to see if some other transmitter is currently active. Such a situation leads to ‘collisions’ where each competing transmitter waits a progressively larger time before attempting retransmission. This characteristic of the original Ethernet made it a poor choice for many types of ‘real time’ communication.

Since all receivers are going to see exactly the same message, even though the message is typically only intended to be received by a single receiver, each message has at its front a pair of 48-bit address fields known as ‘Media Access Control’ fields, also known as ‘MAC addresses’. The first such address indicates the intended target of the message. The second such field indicates the sender of the message.

A message received by a station, but whose target MAC address does not match the MAC of that station, is typically discarded. Usually this decision is made by the Ethernet interface hardware itself, without placing any burden on the computer software, since the matching of two 48-bit numbers is a simple operation to perform using electronics.

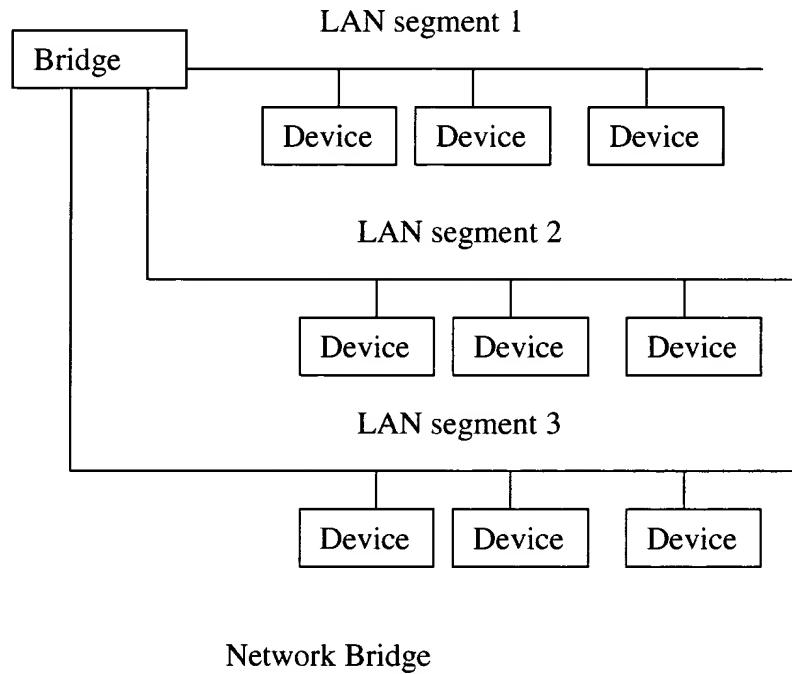
Destination MAC	Source MAC	Message Contents Eg BOOTP, ARP, TCP, whatever	Error Check
-----------------	------------	--	-------------

Layer 2 (eg Ethernet) datagram structure

These MAC addresses are treated by all networking components as having no defined structure or inherent meaning. They are simply assumed to be ‘globally unique’, meaning that no two network interfaces which could ever be attached together on the same broadcast bus would ever be allocated the same 48-bit number. Most manufacturers treat the MAC address of a device very much like a device serial number – it is a number which the manufacturer is free to allocate, out of a range of addresses he has previously obtained from the primary registration authority. This registration authority, in turn, is the Institute of Electrical and Electronic Engineers, or IEEE.

When the original Ethernet specification was being put together, it was recognized that there would always be a need to interconnect Ethernet segments, in order to increase the number of potential devices which could inter-communicate using the Ethernet local area network. So a device was defined, called a ‘Network Bridge’. These devices were in common use by the early 1980’s.

A network bridge is a computer with two or more network interfaces, such as Ethernet interfaces, where each network is attached to a different local area network. In the definitions, these interfaces are referred to as ‘ports’ and usually have port numbers numbered from 1 to N. The receiver portion of each interface ‘listens’ to the traffic on each interface, and attempts to determine whether the sender and receiver happen to be on the same ‘side’ of the bridge, or on different ‘sides’.



If the bridge determines that a message has been received on a particular network port, and the bridge believes that the target of the message is accessed through a different port, then the bridge will take responsibility for ‘storing and forwarding’ the message. To do this the bridge accepts the message completely, and then as soon as possible, it retransmits the message verbatim on the ‘outbound’ port which it believes is appropriate to reach the target in question. If, on the other hand, the target is believed to be on the same side as the message was received, the bridge will deliberately refrain from copying the packet.

The effect is that messages only pass through those portions of the extended network which are absolutely necessary to reach their target. The messages are not propagated into any network segments unrelated to the target, with the result that congestion is reduced on those unrelated

segments, and more aggregate traffic can be carried in unit time by the extended network than if the bridges were not present.

The end result of this is to significantly increase the ability of the network to ‘scale’ – that is, to add more and more devices to the same network without increasing its load and response time.

In order for a bridge to accurately deduce on which port to transmit each individual packet, it must first ‘learn’ where each communicating device is. In particular, it must learn on which of its ports messages from that MAC address have been received in the recent past. To do this, the bridge listens for all traffic, and in addition to inspecting the ‘destination MAC’ field, which it would need in order to decide whether to copy the message, it also records the ‘source MAC address’ field.

Now the switch maintains a data table wherein it records the port number associated with each MAC address based upon recognizing said MAC address as being in the ‘Source MAC’ field of a message, and the port number being the port on which the message was received. The table is initially empty, when the bridge is reset or goes through a power cycle. When the first messages are received after such reset, there will be no entries in the data table giving guidance as to the port number to which the message should be forwarded. So the bridge resorts to ‘flooding’ the message, transmitting it on all ports other than the one on which it was received. It does, meanwhile, record the source MAC in its table.

In general, for most messages sent in one direction from a source to a target, there will be a response sent shortly thereafter from the target back to the original source. This message will of course include its own ‘source MAC’ and ‘destination MAC’ fields. So typically, once the response message has propagated back through the bridge, the bridge will have updated its tables to accurately identify the port number for both source and target. Thus flooding of the network, the unnecessary copying of messages to network segments ‘just in case’, will be limited in most cases to a single flooded message per target MAC per reset. And since the number of messages sent between such resets will typically be measured in billions, this reduces to a negligible proportion.

A typical bridge today can keep track of 16,000 or more such MAC addresses, meaning that it is feasible to attach up to 16,000 devices to a single local area network segment without resorting to ‘network flooding’.

Network Switch

Although originally Ethernet segments were based on shared access to a common coaxial cable, the more modern variants use a single twisted copper pair or optical fiber to connect a single transmitter to a single receiver, and vice versa.

Since there is only a single transmitter attached to each cabling segment, there is no longer any need to delay transmission waiting for the channel to become clear. This results in significant improvements in responsiveness, and means that Ethernet networks can now carry traffic where any transport delays would be problematic, such as telephone traffic.

These modern cabling arrangements use what is known as ‘point to point’ wiring, with the transmit and receive cables being combined together in a single sheath, and connected at one end to a port of a bridge, and at the other end to the interface of the target computer or other device.

In this style of wiring, the bridge devices typically have a large number of ports (12-48 being most common), and since the wiring is no longer shared between devices, there is no reason for a device or bridge ever to hesitate before transmission (in case someone else were occupying the channel).

Because of this, the descriptive term used for the unit which performs the bridging function has changed in common parlance to ‘Network Switch’ instead of ‘Network Bridge’. However, the method of operation is still exactly the same, as are the specifications which determine how the bridge/switch must operate.

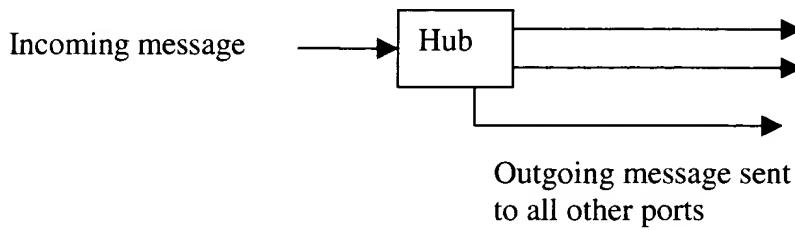
Managed Network Switch

A managed switch is a network switch which is designed to report its internal data and statistics to a management station, using the standard reporting protocol known as ‘Simple Network Management Protocol’ or SNMP.

In particular, the standard document which defines how this reporting is done is known as the ‘Bridge MIB’ or ‘RFC 1493’. Any network switch which claims in its advertising or data sheet to support RFC 1493 is known in the networking community as a ‘Managed Network Switch’.

Network Hub

A network ‘hub’ or ‘Ethernet hub’ is a simple amplifying device which causes any Ethernet message received on one of its ports to be copied in identical form to every other one of its ports. It is a much less complex (and expensive) device than a Bridge, and was for many years the standard method of adding additional devices to a single Ethernet local area network.



Use of hub

By its nature, a hub has no interest in recording the identity of the port on which a particular message was received. The vast majority of Ethernet hubs installed today provide no management data via SNMP.

Router

A router, also known today as a ‘Layer 3 switch’ is a device which is designed to connect together multiple local area networks to construct an Internet. The design of routers indeed predates Ethernet

networks, being originally designed as the interconnect device in the ARPANET developed by Bolt Baranek and Neumann, and later adopted as a standard component of the TCP/IP protocol set which is today ubiquitous.

By the very nature of the TCP/IP protocol, the individual local area networks which are interconnected by the routers must have distinct and identifiable ranges of Internet Addresses, or IP subnet addresses.

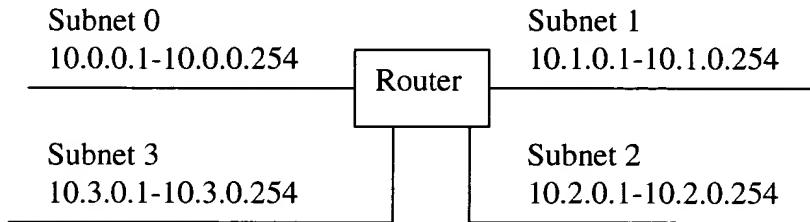
This is a key difference between bridges and routers. In an internetwork connected by routers, known as a ‘Layer 3 network’, the individual networks must have been previously allocated distinct ranges of IP addresses. It is not possible to ‘move’ a device from one subnetwork to another without reconfiguration.

By contrast, an internetwork connected by bridges, known as a ‘Layer 2 network’, all of the devices on all ports of the bridge are considered to be part of the same IP subnet, and there are no restrictions on address assignment. It is perfectly normal for a device previously connected to one port of a bridge, to be disconnected and reattached to a different port without requiring reconfiguration. The bridge is unaware of any associations of devices with IP address ranges.

A message sent to a target whose address is recognized as being within the same IP subnet range is considered a ‘local’ message, and the local area network techniques used for local message delivery would be used. In the case of Ethernet networks, this mechanism involves use of the ‘Address Resolution protocol’, to find the MAC address associated with an IP address, followed by direct delivery of the message using an Ethernet datagram.

A message sent to a target whose address is not local, by contrast, is passed to a router. The router typically has a number of interfaces attached to different known IP subnet address ranges, and maintains a table known as a ‘routing table’ to decide the best route to forward the message to get it ‘closer’ to its destination. In almost all cases, one of the interfaces would be a path to another router known as a ‘default gateway’ to which all messages for unknown IP addresses would be forwarded.

This is the fundamental way in which the Internet operates, each router in the network deferring to a superior one with more knowledge. Typically every major Internet carrier maintains one or more Core routers with tables elaborate enough to forward any possible packet to the correct destination.



Use of Router

Maintenance of the routing tables themselves is an intricate task, normally left to skilled professionals, and assisted by certain routing protocols such as OSPF and RIP. When these routing tables get adjusted incorrectly, it typically causes major disruption to the intranet or Internet.

A router is also almost always a ‘managed device’, meaning that its information can be reported and accessed using SNMP. However a router has no information of value when attempting to deduce the correct unique IP address for a previously unknown device.

The only assistance a router can give is to identify on which interface a BOOTP request was received, and this can be used in many DHCP arrangements to select the correct ‘pool’ of available addresses to pick a random member. This is important because by the nature of TCP/IP, a station which was given an IP address inconsistent with its peers on the same local area network would be unable to communicate.

However, in the case of the present invention, such information would be unhelpful. The IP address which must be allocated has to be exact, not one randomly picked from a pool, and for the routing data to help, there would have to be a DHCP pool size of exactly one for each subnetwork. An internetwork where each subnetwork had only a single member would be unusable (the router itself which communicates with the subnet consumes already a single IP address)

Significance of Managed Network Switch to Present invention

The problem which the present invention solves is the ability to connect a device, which has previously not been configured in any way other than being preloaded with software at a factory, and have this device acquire the correct IP address by arranging that it is attached using the same network cable as a failed predecessor and processing as noted in the patent application.

In order to achieve this, it is necessary for the mechanism to augment the conventional BOOTP server operation, which assigns IP addresses based upon previous (manual) configuration of the MAC addresses of the devices, by having the BOOTP server adjust its own database automatically based on SNMP position reporting.

The BOOTP request message is itself, of course, an Ethernet datagram. It contains a source MAC address, which would ordinarily be inspected by the BOOTP server directly to identify the correct IP address. But in the case of a replacement, the MAC address will in fact not be already known to the BOOTP server, since the device has not been used previously on the network.

However, in order for the BOOTP message to reach the BOOTP server, it must have passed through one or more managed network switches. Each of these switches will have noticed the ‘source MAC address’ of the message, and the port of the switch on which the message was presented. The data will have been recorded in the bridging tables of the switch, and in the case of a managed switch will be made available for inspection by a management station using the SNMP protocol.

The BOOTP server, having noticed that the MAC address is one which is NOT in its table, is enhanced to take on the role of such a management station, and makes a series of SNMP queries to the managed network switches on the network.

Each such query asks whether the switch concerned has recorded the port number on which was received a message from the designated source MAC address. Each switch which has such information will provide that in the response which it sends to the BOOTP server. In turn, the

BOOTP server determines from the set of such position reports whether the combination of managed network switch address, and port number of that managed network switch, matches exactly with the equivalent coordinates of a device which has previously failed.

If the appropriate conditions are satisfied, the BOOTP server updates its table to associate the IP address with the newly-seen MAC address, and issues the BOOTP response to the requesting target device. All of this happens typically within less than a second after the replacement device is first attached to the network, and without any human intervention.

Nothing in the Tosey patent would guide one experienced in the art to combine the capabilities of recording the port number of a managed network switch on which a message from a previously unknown station was received; and requesting that information using the SNMP protocol; and altering the IP-MAC configuration table for a BOOTP server as a result of such interrogation. Nor is there any reason for Tosey to arrive at the present invention because it is intended to solve a different problem.

The undersigned declares that all statements of his own knowledge made herein are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application of any patent issuing thereon.

Respectfully submitted,



Andrew G. Swales

21 July 2005
Date

Applicant's Attorneys:
Scott J. Asmus, Reg. No. 42,269
Maine & Asmus



ANNOTATED SHEET

2/8

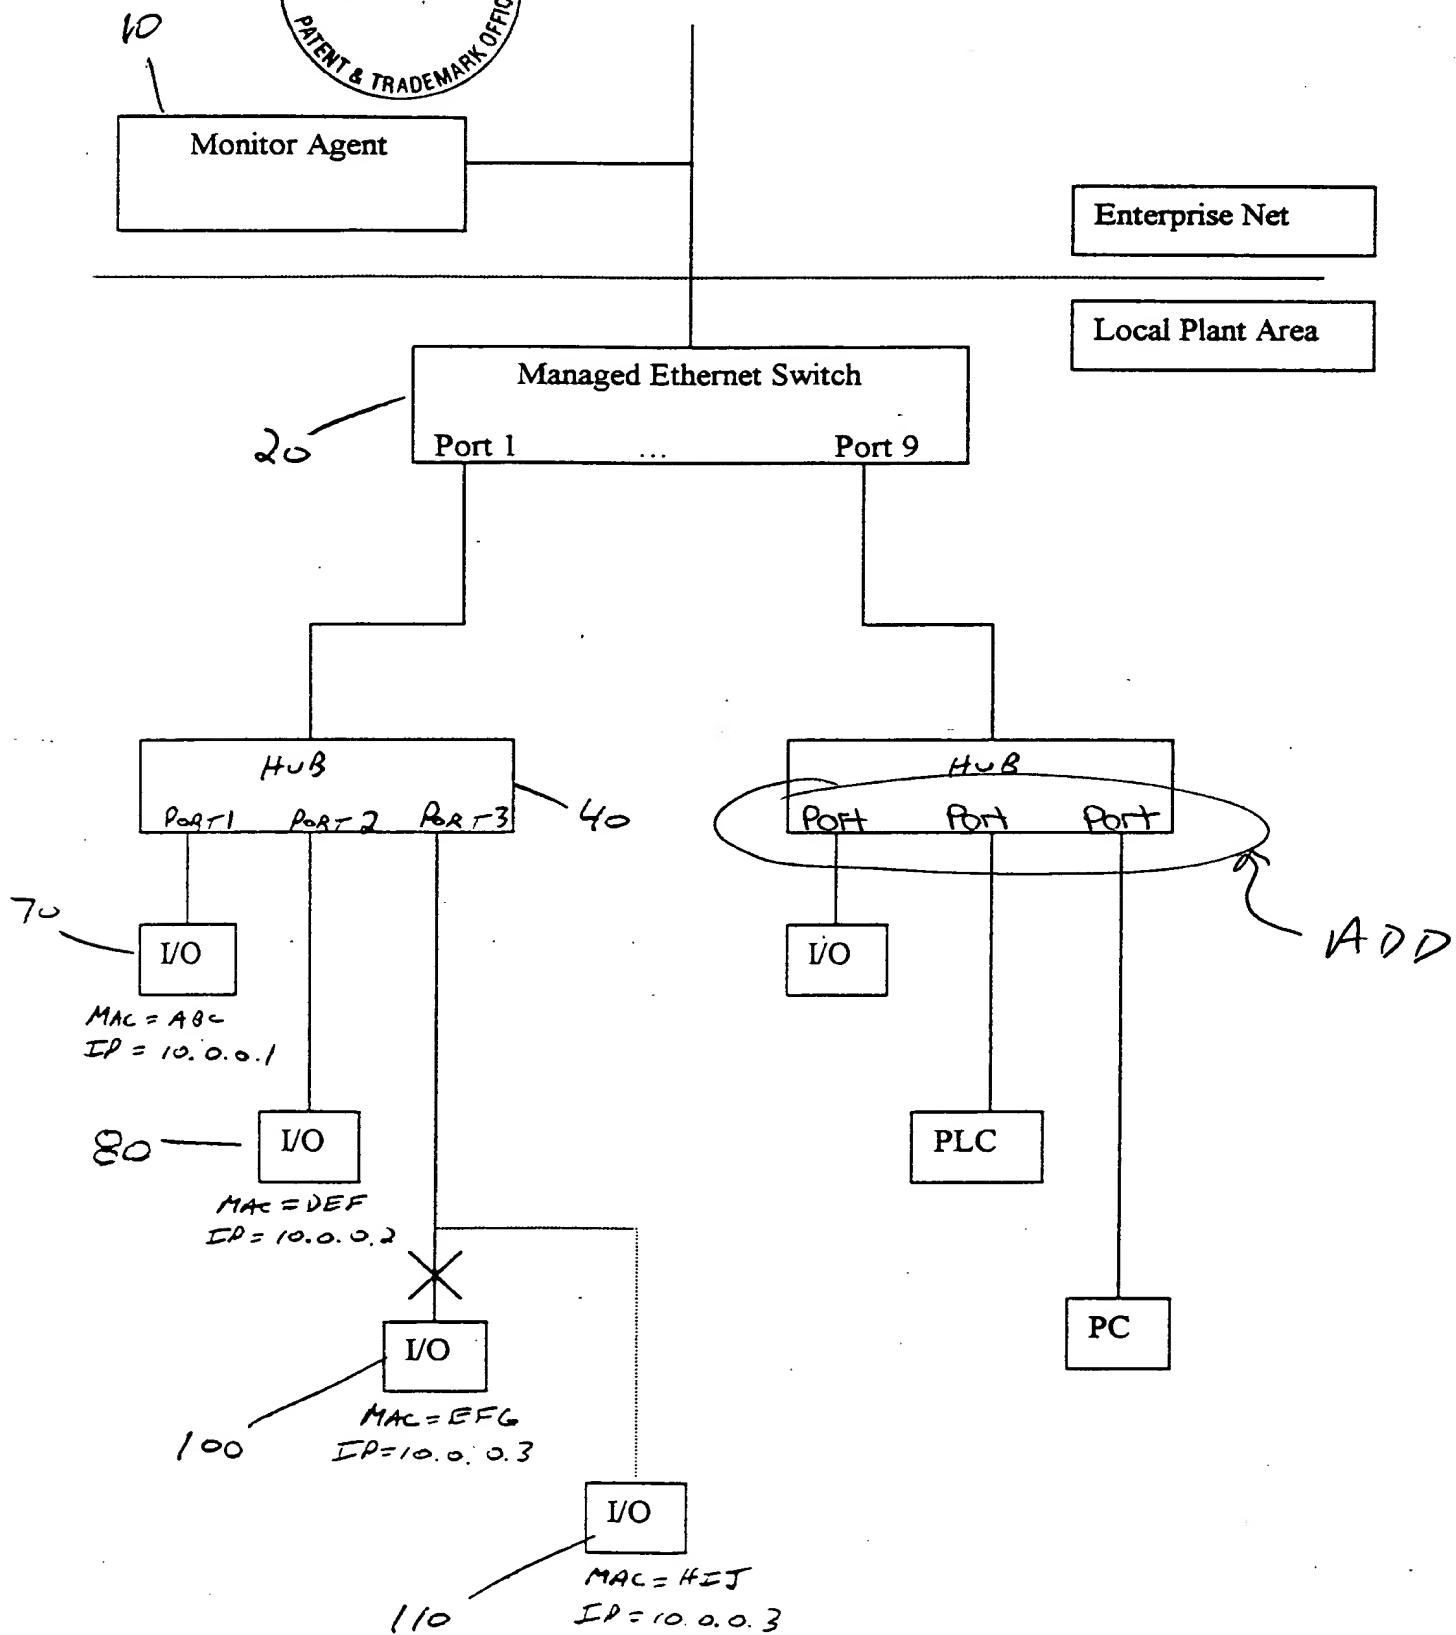
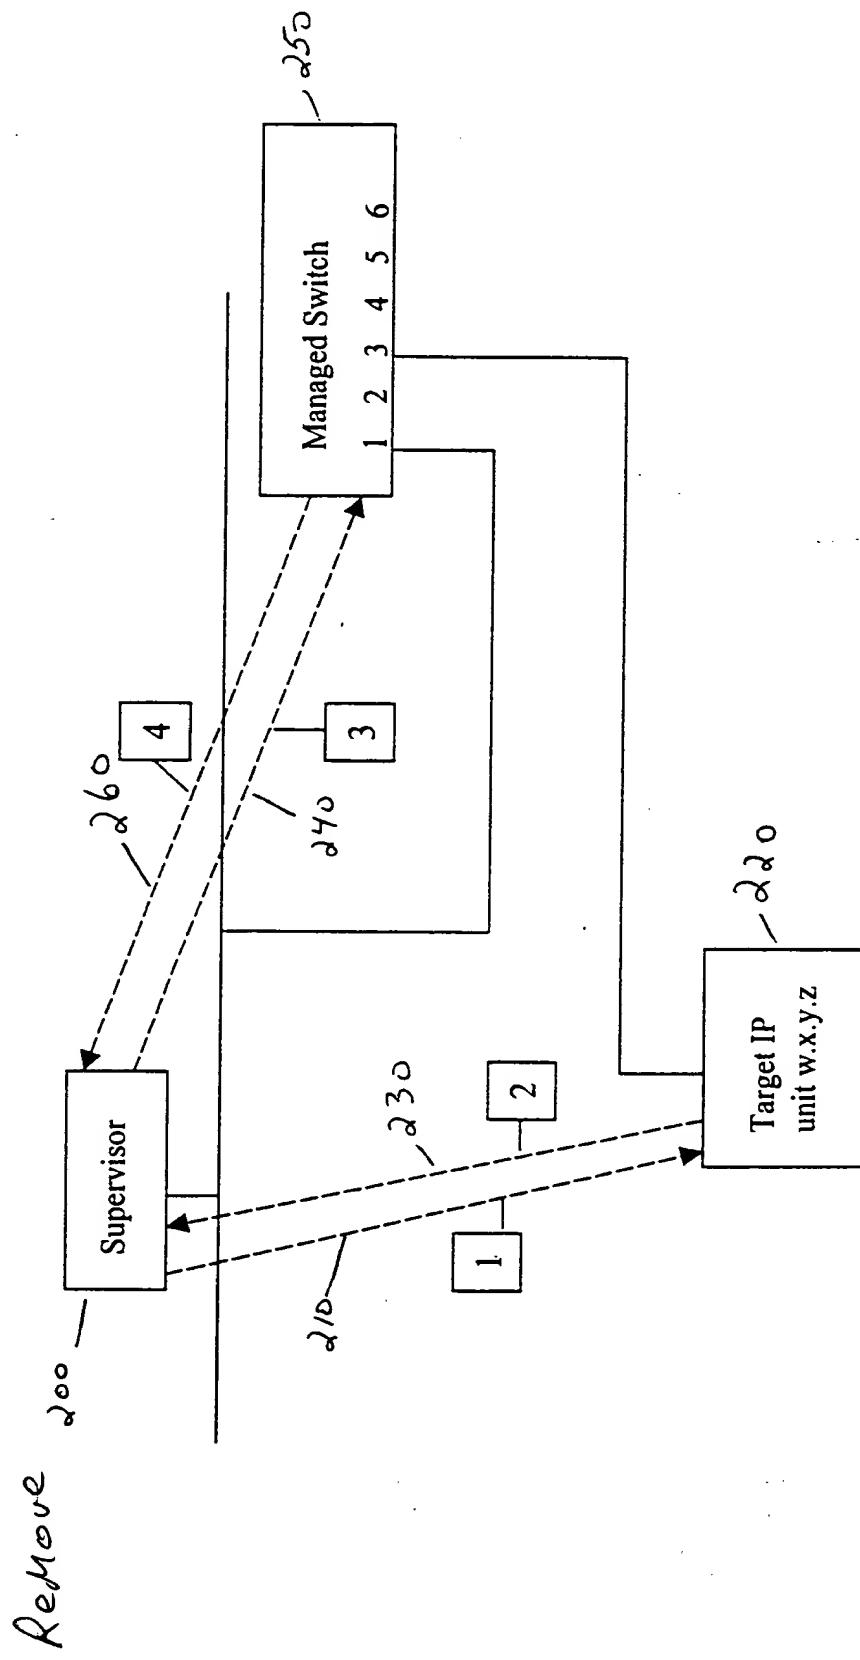


FIG 2

ANNOTATED SHEET

3/8

Discovery=/confirmation=sequence. Find MAC address of target and record its canonical location (numbered port-of-supervised-switch) Dedicated port scenario

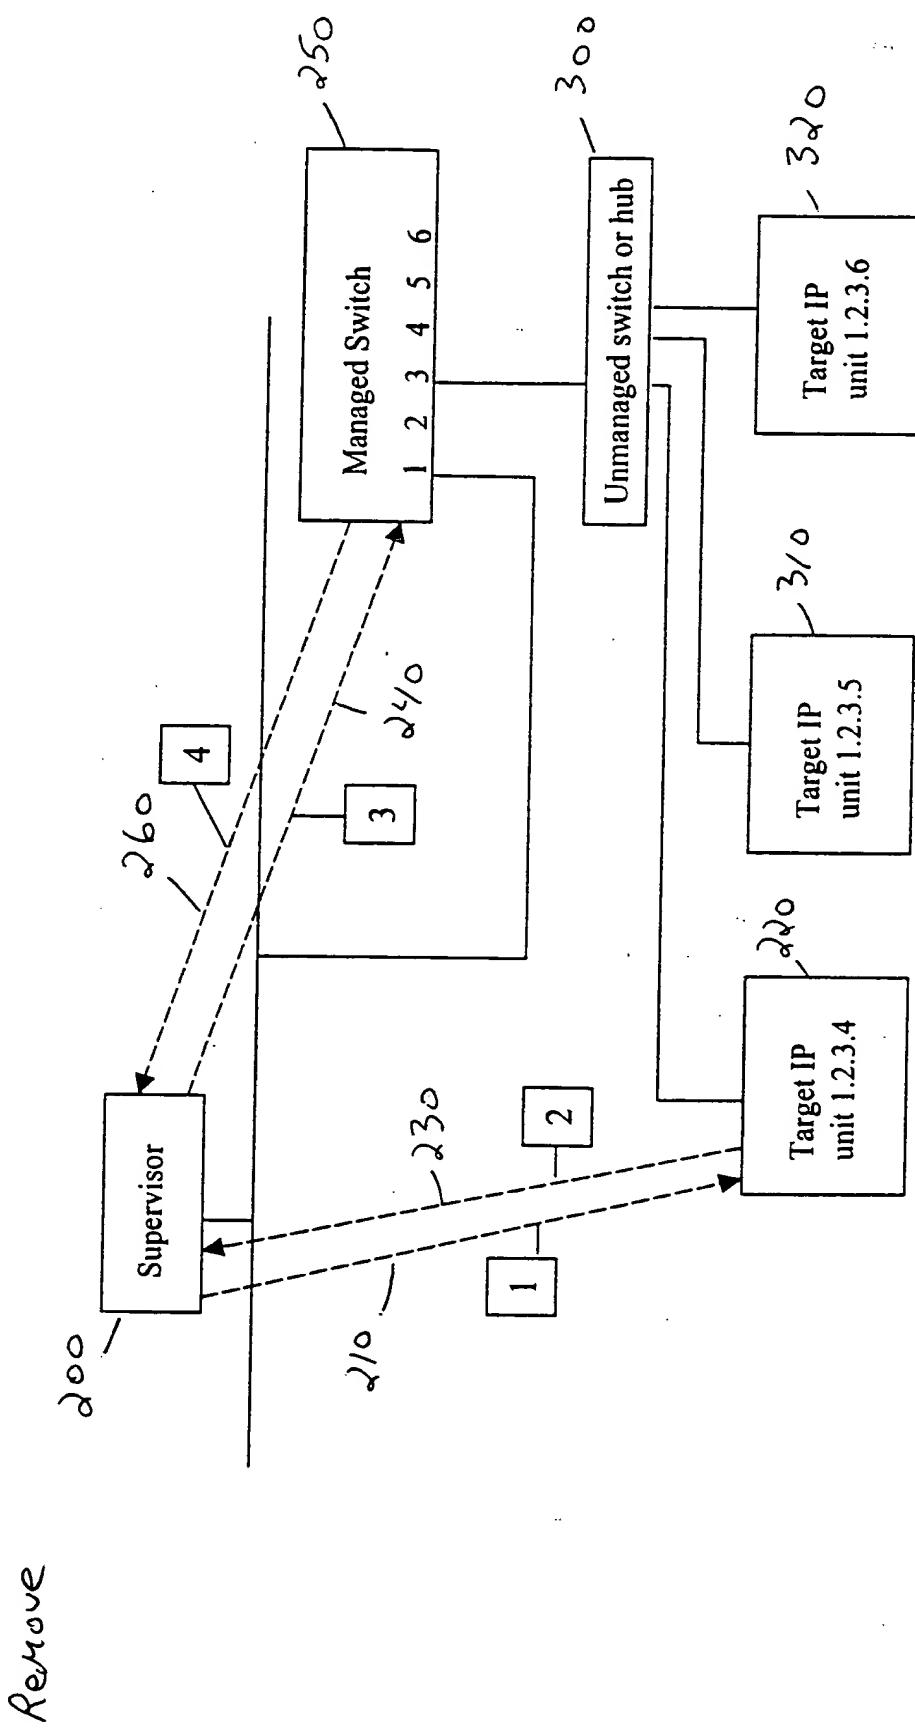


1. ARP Request - inquire MAC address of IP address w.x.y.z (broadcast)
2. ARP response - MAC address of requested IP address is xxx
3. SNMP Findport request - request port number of MAC xxx
4. SNMP Findport response - port number of MAC xxx was 3

FIG 3

ANNOTATED SHEET

4/8



1. ARP Request - inquire MAC address of selected IP address 1.2.3.4 ~~(68.40.24.57)~~
2. ARP response - MAC address of requested IP address is xxx
3. SNMP Findport request - request port number of MAC xxx
4. SNMP Findport response - port number of MAC xxx was 3

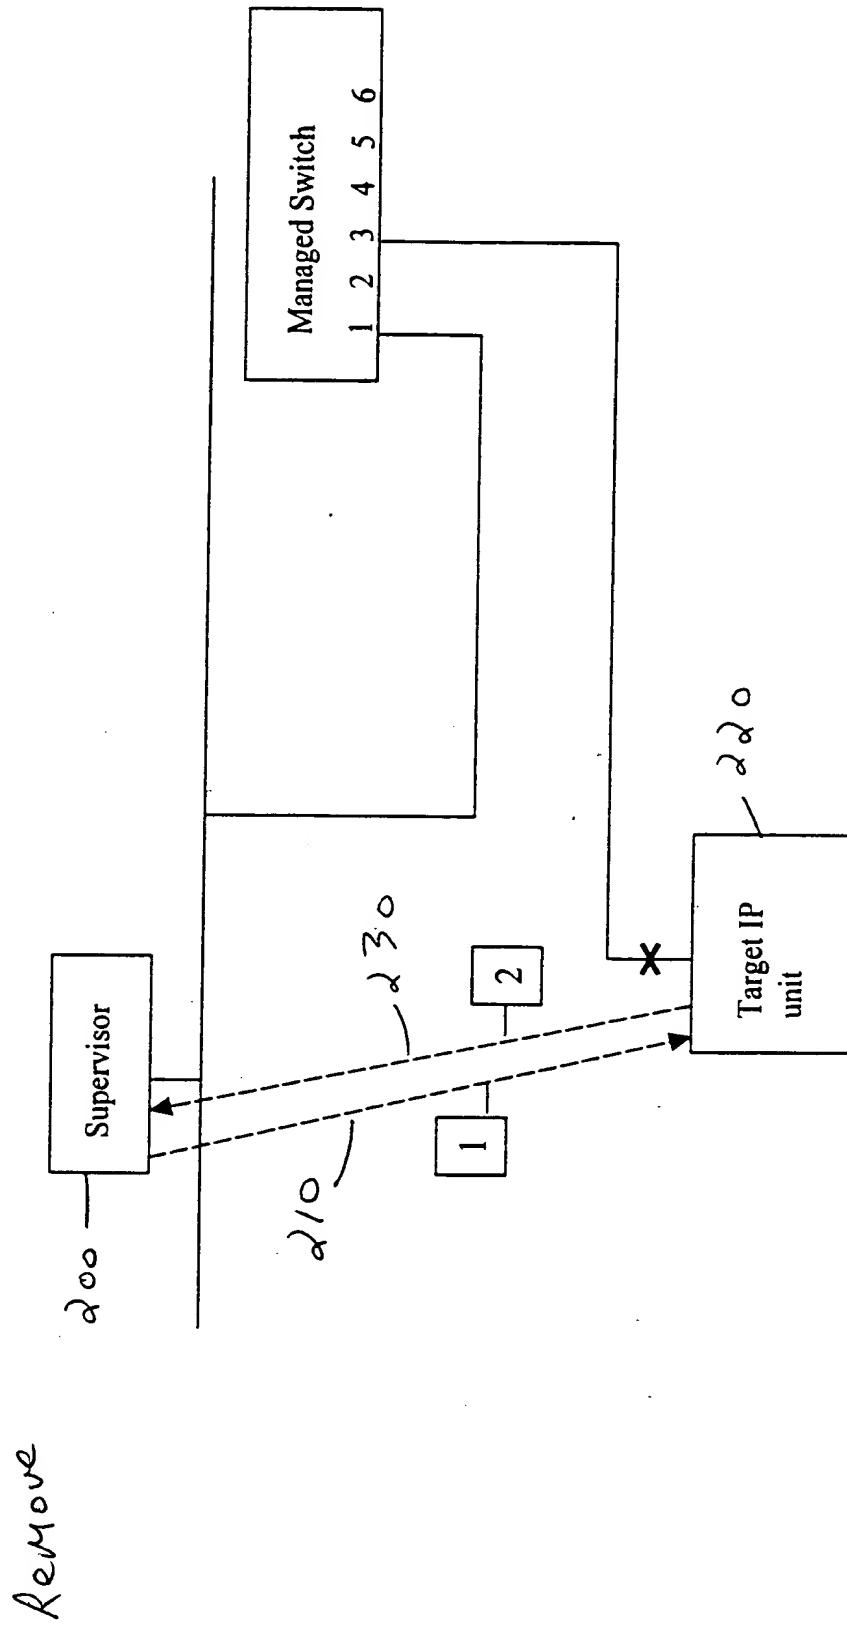
Targets are automatically determined to be sharing port 3 of the switch.

FIG 4

ANNOTATED SHEET

5/8

Confirm-presence-sequence - interrogate-target at periodic interval. Absence-of response indicates target 'down'. A single target down in a canonical location becomes-a-reassignment-candidate. Dedicated port scenario.

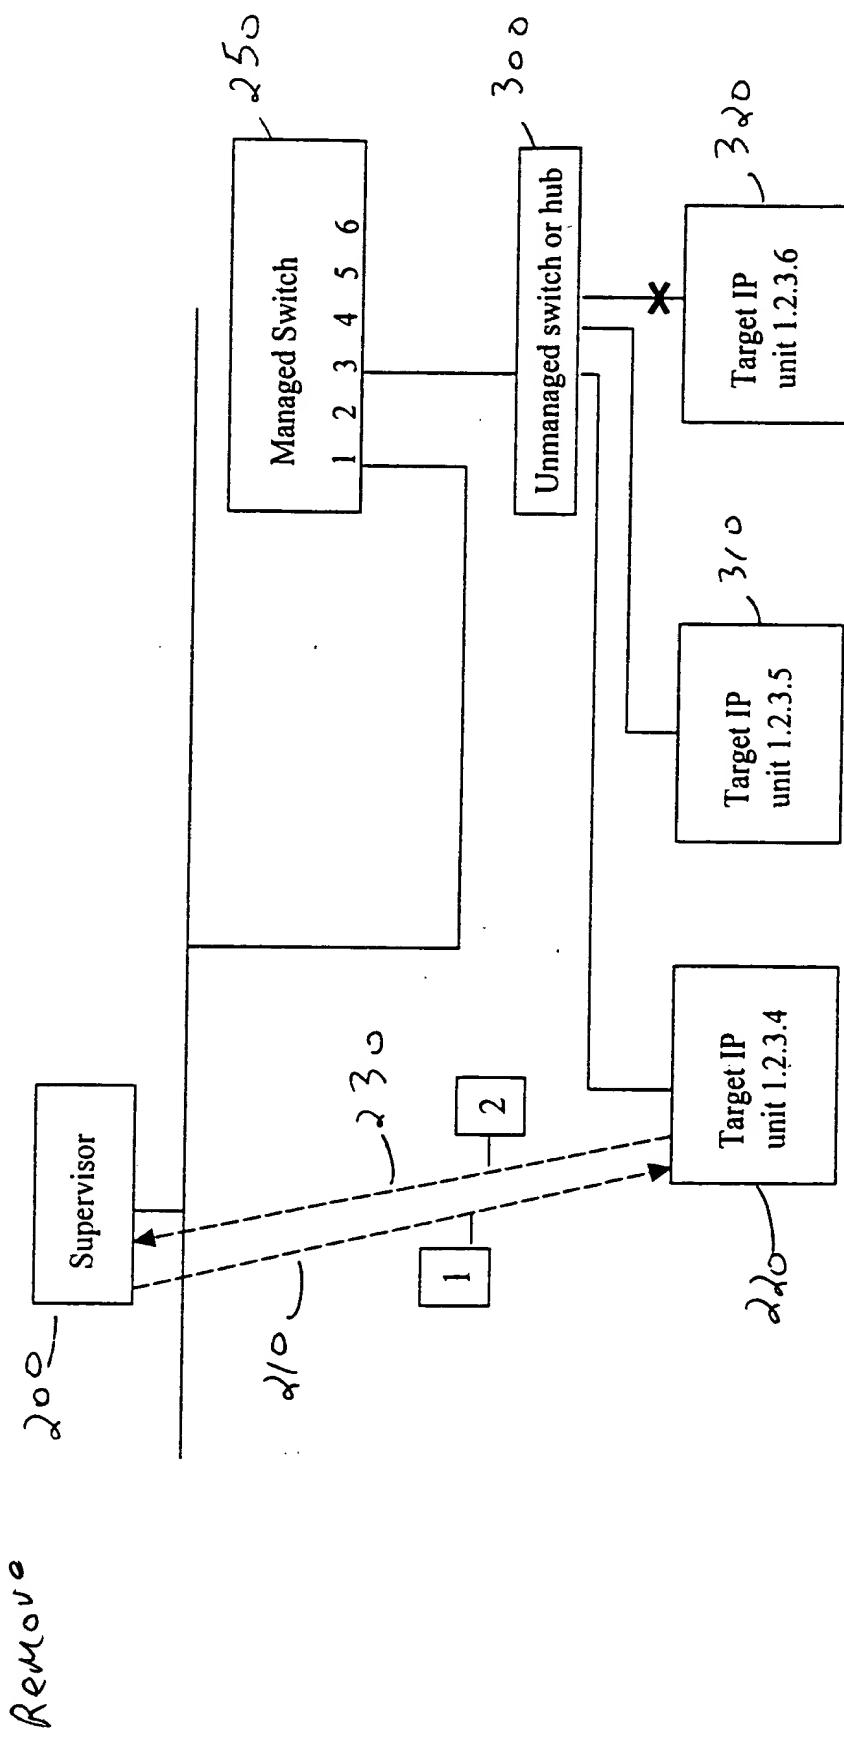


1. ARP Request - inquire MAC address of selected IP address (unicast)
2. ARP response - MAC address of requested IP address is xxx
If no response is received, signify that the target IP unit is 'down'

FIG 5

ANNOTATED SHEET

6/8



1. ARP Request - inquire MAC address of selected IP address 1.2.3.6 (unicast)
2. ARP response - MAC address of requested IP address is xxx
If no response is received, signify that the target IP unit is 'down'
If 1.2.3.6 is the ONLY unit on port 3 of the switch which is down, then it is a candidate for automatic reallocation. If any of the other units found on this port are also down, automatic reallocation will be suppressed.

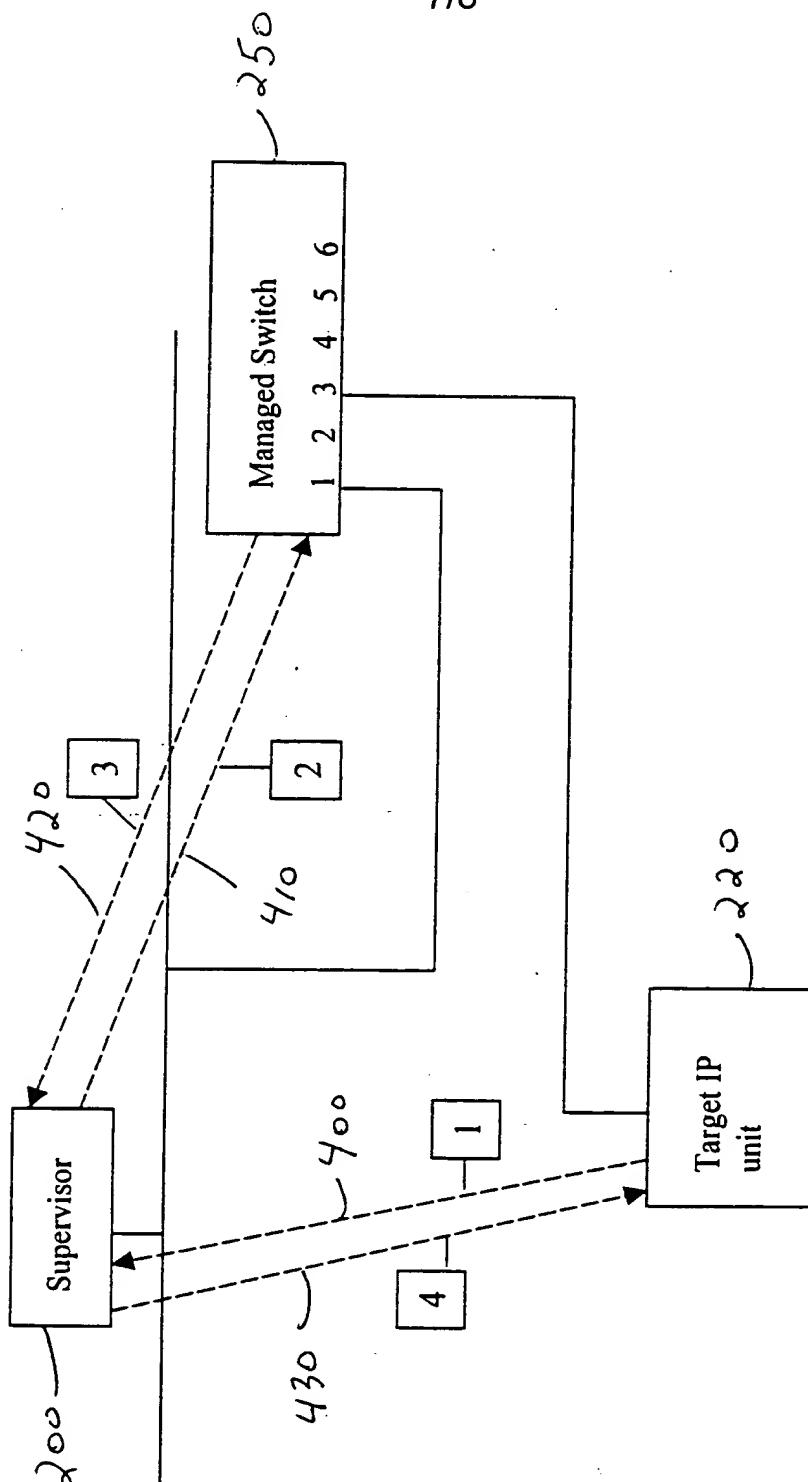
FIG 6

ANNOTATED SHEET

7/8

IP address assignment sequence. Target IP broadcasts request for address. The target was previously running at that location (eg. just reset or power cycled).

Remove



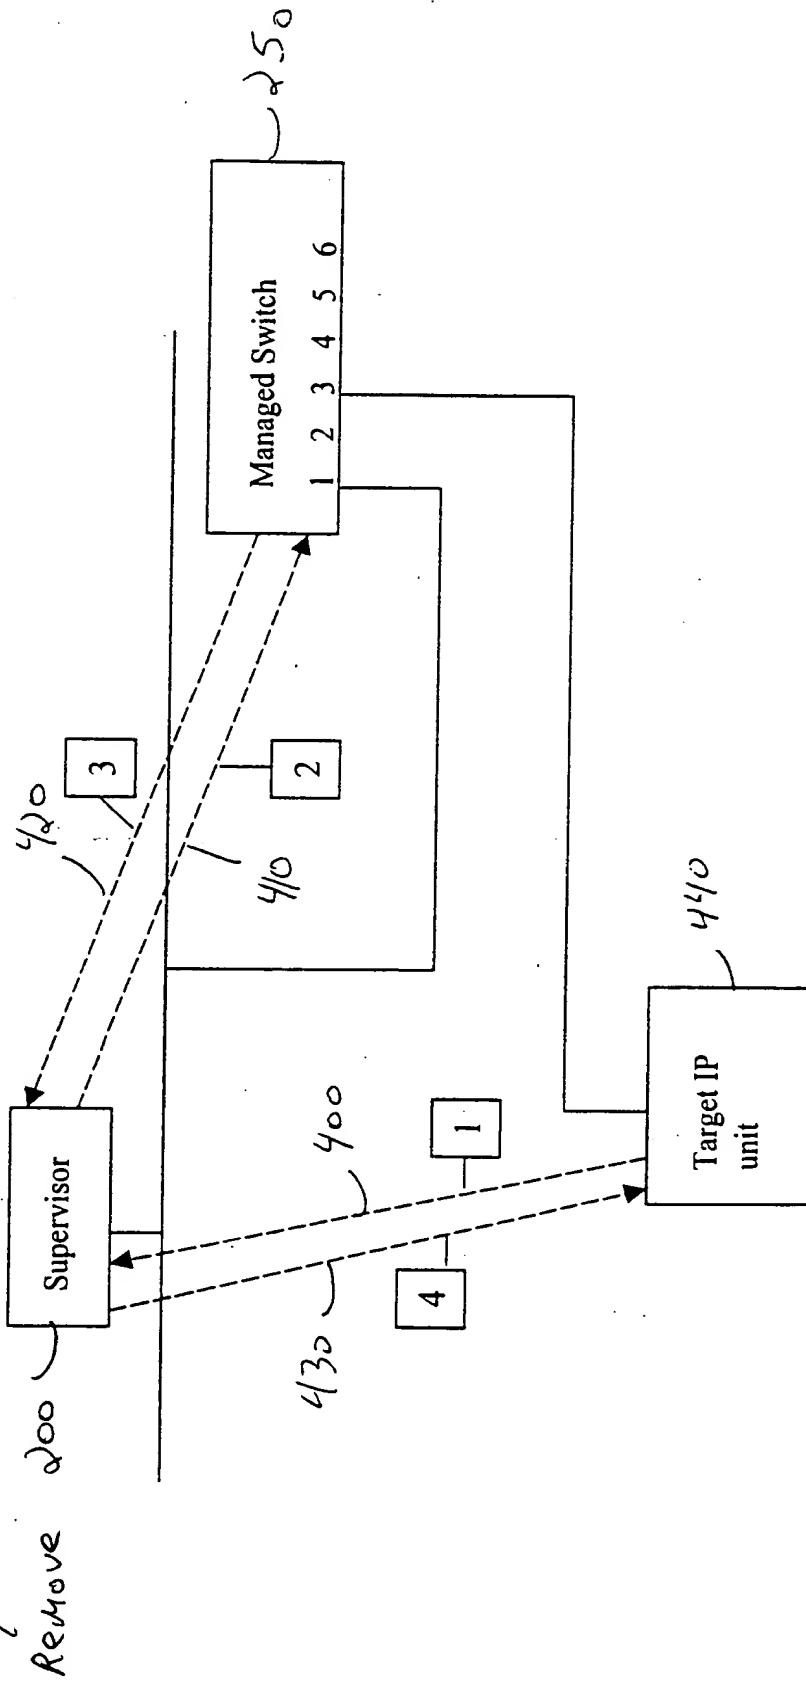
1. BOOTP request - please supply IP address for MAC xxx (broadcast)
2. SNMP Findport request - request port number of MAC xxx
3. SNMP Findport response - port number of MAC xxx was 3
(MAC xxx already associated with IP at that canonical location - OK to assign)
4. BOOTP response - IP address for MAC xxx is w.x.y.z

FIG 7

ANNOTATED SHEET

8/8

IP address-reassignment sequence. Target IP broadcasts request for address. The target was NOT previously running at that location. A single target IP unit at that location is determined to be currently 'down', and is assumed to have been replaced with another using the same cable.



1. BOOTP request - please supply IP address for MAC xxx (broadcast)
2. SNMP Findport request - request port number of MAC xxx
3. SNMP Findport response - port number of MAC xxx was 3
(MAC xxx not known. However a single unit at that location is not currently responding. Update BOOTP equivalence table to record new IP assignment and authorize assignment)
4. BOOTP response - IP address for MAC xxx is w.x.y.z

FIG 8